

Automated Pynta-Based Curriculum for ML-Accelerated Calculation of Transition States

Trevor Price, Saurabh Sivakumar, Matthew S. Johnson, Judit Zádor,* and Ambarish Kulkarni*



Cite This: *J. Phys. Chem. C* 2025, 129, 7751–7761



Read Online

ACCESS |



Metrics & More




Article Recommendations



Supporting Information

ABSTRACT: Microkinetic models (MKMs) are widely used within the computational heterogeneous catalysis community to investigate complex reaction mechanisms, to rationalize experimental trends, and to accelerate the rational design of novel catalysts. However, constructing these models requires computationally expensive and manually tedious density functional theory (DFT) calculations for identifying transition states for each elementary reaction within the MKM. To address these challenges, we demonstrate a novel protocol that uses the open-source kinetics workflow tool Pynta to automate the iterative training of a reactive machine learning potential (rMLP). Specifically, using the silver-catalyzed partial oxidation of methanol as a prototypical example, we first demonstrate our workflow by training an rMLP to accelerate the parallel calculation of DFT-quality transition states for all 53 reactions, achieving a 7× speedup compared to a DFT-only strategy. Detailed analysis of our training curriculum reveals the shortcomings of using an adaptive sampling scheme with a single rMLP model to describe all reactions within the MKM simultaneously. We show that these limitations can be overcome using a balanced “reaction class” approach that uses multiple rMLP models, each describing a single class of similar transition states. Finally, we demonstrate that our Pynta-based workflow is also compatible with large pretrained foundational models. For example, by fine-tuning a top-performing graph neural network potential trained on the OC20 dataset, we observe an impressive 20× speedup with an 89% success rate in identifying transition states. This work highlights the synergistic potential of integrating automated tools with machine learning to advance catalysis research.



INTRODUCTION

Microkinetic models (MKMs) derived from density functional theory (DFT) are a powerful tool in the field of heterogeneous catalysis,^{1,2} linking DFT-predicted kinetics to experimental observables such as reactant conversions and product selectivities.^{3,4} These models have been successfully leveraged for optimizing reactor conditions, elucidating rate-determining steps, and high-throughput catalyst screening.⁵ However, as advanced experiments provide more detailed insights into these complex systems, such as site-specific kinetics⁶ and the coupling of surface and gas-phase reactions,^{7–17} there is a growing demand for models that can better capture the complexity of real catalytic processes. Traditional approaches in the construction of DFT/MKMs are manually intensive and computationally expensive, limiting their comprehensiveness and necessitating the use of approximations which reduces their accuracy.¹⁸ Coupling automation with strategies to reduce computational expense provides a pathway to address these challenges and accelerate MKM development.

Automated tools such as Reaction Mechanism Generator (RMG)^{19–24} facilitate rapid generation of coupled heterogeneous and homogeneous MKMs.^{3,25,26} RMG leverages internal databases to estimate unknown thermochemistry and kinetic parameters. However, for surface-mediated reactions, these

databases are sparsely populated and often rely on approximations like Brønsted–Evans–Polanyi (BEP) relationships²⁷ to estimate reaction barriers. These approximations often carry large errors (e.g., >0.2 eV²⁸) and thus reduce the predictive capacity of the model. Additionally, BEP relationships do not provide actual transition state geometries and are typically used for investigating activity trends across different catalysts.

Obtaining accurate reaction barriers requires DFT calculations of transition state energies, wherein the most difficult task is to locate the corresponding first-order saddle point. The nudged elastic band (NEB) method is commonly used for this, but it is computationally expensive as it requires optimizing a series of images along the reaction coordinate.^{29,30} This expense often leads to neglecting the possibility of multiple adsorbate or transition state (TS) structures for a given reaction, making it difficult to unambiguously identify the

Received: January 14, 2025

Revised: March 6, 2025

Accepted: March 24, 2025

Published: April 13, 2025



lowest-energy saddle point. Given the exponential dependence of the rate coefficient on the DFT-derived barrier height, this can result in rate coefficients that are orders of magnitude too high. New approaches are needed to lower computational costs and avoid approximations which limit the application of MKMs.¹⁸

Machine learning potentials (MLPs) can accelerate the exploration of the potential energy surface by predicting forces at a fraction of the DFT cost, with accuracy approaching DFT when trained on sufficiently large DFT databases.³¹ Two promising strategies to speed up saddle-point searches with MLPs are (1) active learning to create system-specific models and (2) using large pretrained models. Active learning has significantly sped up NEB calculations using Behler–Parrinello neural networks and Gaussian process regression models but has not yet been applied to large reaction sets needed for MKMs.^{32–35} Another approach demonstrated active learning of a Gaussian approximation potential (GAP)³⁶ model to accelerate NEB calculations and to run dynamic studies to bypass the common harmonic approximation. However, GAP models costs can scale unfavorably with system size, limiting its use for large reaction networks.^{37,38}

Large DFT datasets like OC20³⁹ have enabled the development of general and accurate pretrained models.⁴⁰ Wander et al.⁴¹ introduced CatTSunami, a framework which leverages OC20-trained³⁹ models to accelerate comprehensive TS exploration using NEBs, achieving impressive speedups across various surfaces and reaction types. However, the performance declines for more complex reactions (e.g., desorptions versus dissociations and transfer reactions). Follow-up work demonstrated that using the Hessian-based optimizer Sella^{42–44} improves the convergence of transition states explored with CatTSunami, though some reactions still pose challenges.⁴⁵ While promising, pretrained models require fine-tuning for broader use in systems studied with different DFT parameters compared to those used in generating the model's training dataset. While fine-tuning has been demonstrated on adsorbates,^{46,47} this has not yet been demonstrated for transition state optimization. Larger models also carry much higher computational cost per inference, which can limit their application in dynamical studies to directly obtain free energy barriers beyond the harmonic approximation.³⁶

This work demonstrates an automated framework that can leverage both smaller models (trained through active learning) and large pretrained models (via fine-tuning) for accelerating transition state optimization. Currently, to the best of our knowledge, there does not exist an automated, general, and parallelized workflow for training system-specific models for accelerating transition state searches across the large sets of reactions required for MKM development. In this work, this was achieved by integrating Pynta⁴⁸ (<https://github.com/zadorlab/pynta>), an automated kinetics workflow tool for surface reactions, within a training curriculum for machine learning potentials based on neural network architectures. Specifically, we leverage Pynta's comprehensive transition state exploration to automatically sample the potential energy surface along the reaction coordinate regions for several reactions in parallel. Here, the reaction coordinate region refers to the configuration space explored by Pynta during our workflow. This is achieved by generating many transition state guess structures, their optimization to a first-order saddle point (via Hessian-based optimization), and confirmation that the resulting optimized transition states connect the intended

reactants and the products (via intrinsic reaction coordinate optimization). We apply commonly used ensemble-based techniques⁴⁹ to quantify model uncertainty, facilitating sampling of a set of the Pynta-generated configurations to iteratively construct a DFT dataset for training an rMLP. Tools for automatically determining valid transition states and generating the thermochemical and kinetic databases for updating RMG are now provided within the Pynta workflow.

The effectiveness of this Pynta-based workflow is demonstrated on a diverse set of 53 reactions (Table S1) described by Aljama et al.²⁸ to generate an MKM for Ag-catalyzed methanol oxidation, an important industrial process for producing the commodity chemical formaldehyde.^{50–52} We demonstrate a total model approach whereby we train a single model simultaneously on this set of 53 reactions in parallel. Additionally, we also explore a reaction-class-based approach which constructs a model for seven distinct reaction classes (see Table S1) which can take advantage of information sharing, previously demonstrated to improve MLP performance.³⁴ Together, this rMLP-accelerated Pynta approach provides accurate transition state geometries for all 53 reactions at a fraction of the computational cost compared to those of comprehensive DFT/NEB-based approaches. Our analyses show that the Pynta-based curriculum is seven times cheaper than using Pynta/DFT. We recognize that while several computational methods (e.g., RPA and DFT-meta-GGA) are more accurate than DFT-GGA, here, we use DFT-GGA-derived energies and forces for training, as this level of theory is typically employed to describe metal-catalyzed reactions. While beyond the scope of this study, we note that machine learning models have also been trained using beyond-DFT calculations,⁵³ which could be applied within our workflow to improve the accuracy of the derived MKM. Additionally, the ability to obtain transition state geometries at a reduced computational cost can facilitate applying various schemes for correcting DFT energies as demonstrated in the literature.⁵⁴ With the automated generation of thermochemistry and kinetic libraries for updating RMG, this means the reaction network can be easily expanded to include reactions not present in the initial reaction dataset. Our concurrent learning scheme also facilitates using any MLP architecture. Additionally, we demonstrate that a fine-tuned OC20³⁹-trained model can also accelerate the Pynta workflow. This work underscores the importance of the synergistic coupling of automation with MLP acceleration to enhance the accuracy and comprehensiveness of experimentally relevant MKMs.

METHODS

Electronic Structure Theory. DFT calculations were performed using the Vienna ab initio simulation package (VASP).⁵⁵ Electronic cores were represented with projector-augmented wave pseudopotentials.⁵⁶ ASE⁵⁷ was used to interface with VASP. Exchange–correlation was treated with the RPBE functional⁵⁸ with Grimme's D3 dispersion corrections using a Becke–Johnson damping scheme.^{59–61} The bulk lattice constant for Ag used for our slab was determined in our previous study to be 3.891 Å,⁶² in reasonable agreement with the experimental value of 4.079 Å.⁶³ An Ag(111) surface was modeled in this study using a 3 × 3 × 4 slab, where the bottom two layers of the slab were constrained, while the top two layers were allowed to relax. A plane-wave energy cutoff of 500 eV is used. Ionic relaxation steps are terminated when the forces on all of the atoms are

less than 0.05 eV/Å. All slab calculations use a $4 \times 4 \times 1$ Monkhorst–Pack k -point mesh⁶⁴ with >10 Å vacuum spacing between periodic images in the z direction. For slab and adsorbate configurations, electron orbitals were smeared with a first-order Methfessel–Paxton method with a smearing width of 0.2 eV. Gas-phase calculations were performed in a 10 Å^3 box at the gamma point by using Gaussian smearing with a smearing width of 0.05 eV.

Model Parameters and Training. We used the DeepPot-SE^{65,66} within the DeepMD-kit⁶⁷ framework as the MLP architecture. Similar to previous studies within our group,^{62,68,69} we began by using a 3-layer embedding (i.e., 16/32/64) and fitting nets (i.e., 64/64/64). The distance cutoff radius of 6.0 Å with smoothing beginning at 5.5 Å is used. The prefactors for the energy and force contributions to the loss function are $\text{pestart} = 0.02$, $\text{pelimit} = 1$, $\text{pfstart} = 1000$, and $\text{pflimit} = 1$. A standard test–train split of 80% (training), 10% (validation), and 10% (test) is used. We note that the model hyperparameters were not optimized during the concurrent learning workflow. Hyperparameter tuning studies for the final (i.e., iteration 7) model are discussed later in the curriculum limitations and potential improvements sections.

Ensemble-Based Uncertainty Quantification. Sampled configurations are selected for DFT labeling based on their estimated model uncertainties. Similar to previous work by us^{62,68,69} and others,⁷⁰ we use the ϵ_t metric to quantify model uncertainty. Here, ϵ_t for a given configuration R_t is defined by eq 1.

$$\epsilon_t = \max_j \sqrt{\langle \|F_{w,j}(R_t) - \langle F_{w,j}(R_t) \rangle\|^2 \rangle} \quad (1)$$

where $F_{w,j}$ denotes the force on the atom with index j predicted by model w , $\|x\|$ is the norm of quantity x , and $\langle x \rangle$ is the expectation value of quantity x , which is approximated by the average of the ensemble. This metric measures the maximum standard deviation in force predictions obtained using an ensemble of models.⁷⁰ In this work, we use an ensemble of four models that differ in the seed value used to initialize the weights of the neural nets.

Thermochemistry and Kinetics. We included an automated scheme in Pynta to enable the generation of thermochemistry and kinetic libraries for updating RMG. We leverage a prior approach developed by Blondal et al.²⁶ to derive enthalpies of formation of adsorbates using gas-phase reference species with tabulated enthalpies of formation in the Active Thermochemical Tables (ATcT).^{71,72} Vibrational modes are converted to partition functions with the harmonic oscillator approximation. For weakly bound physisorbed species, where calculations show two or more vibrational frequencies $<100 \text{ cm}^{-1}$, low frequency modes correspond to frustrated translation and are replaced with partition functions derived from a 2D gas model. Equations 2–5 were used for the 2D gas where m is the mass of the adsorbate and a is the surface area per binding site of our unit cell (6.6 Å^2). Rate coefficients are calculated based on work from Campbell et al.⁷³ and further described by Johnson et al.⁴⁸ in the original Pynta work.

$$q_{\text{trans}} = \left(\frac{2\pi mk_{\text{B}}T}{h^2} \right) a \quad (2)$$

$$S^{\text{trans}} = R(2 + \log(q_{\text{trans}})) \quad (3)$$

$$\Delta H^{\text{trans}} = RT \quad (4)$$

$$c_p^{\text{trans}} = R \quad (5)$$

RESULTS AND DISCUSSION

Pynta-Based Training Curriculum. Figure 1 illustrates an overview of our three-stage workflow. Stage 1, denoted as

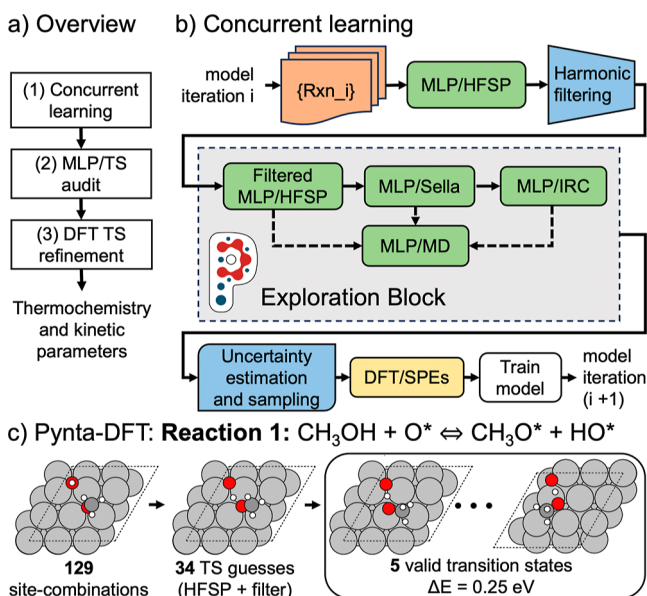


Figure 1. The three stages of our Pynta/MLP workflow are shown in (a). The concurrent learning block (b) uses Pynta to simultaneously explore the reaction coordinate regions for all reactions. Configurations generated in the exploration block (gray box) are evaluated for model uncertainty, downsampled for DFT labeling, and added to the DFT dataset for training the next model iteration. Panel (c) highlights the comprehensiveness of Pynta’s TS search algorithms finding several valid transition states with energy barriers spanning 0.25 eV for the hydrogen transfer reaction from methanol to surface-bound oxygen.

concurrent learning, employs a Pynta-based protocol to explore system configurations along the reaction coordinate region and iteratively train the rMLP. The term “concurrent learning” and its distinguishing features from pure active learning are discussed in prior work.⁷⁰ In short, active learning leverages a set of unlabeled atomic configurations and selects the most relevant data points for labeling. In concurrent learning, there is no data to begin with, and data are generated on the fly as the training proceeds. Concurrent learning can help keep the training dataset small while ensuring the dataset is representative enough for the intended tasks. Here, we focus on obtaining transition states through Pynta’s comprehensive transition state search. Next, the second stage audits these TS configurations to ensure that the TSs obtained in stage 1 correspond to the correct elementary steps. Automation of this step is key, as manually evaluating the set of transition states for each reaction is extremely tedious. Finally, stage 3 refines the audited TSs (from stage 2) with DFT. Taken together, this curriculum provides DFT-quality reaction barriers at a fraction of the computational cost required for Pynta/DFT studies yet preserves the advantages of the comprehensive search aiming to find the lowest-energy adsorbate and TS structures. Note that the results from stage 3 are automatically converted within

Pynta into thermochemistry and kinetic parameters using the approach developed by Blondal et al.²⁶ and can be used to generate databases for inclusion into RMG for microkinetic modeling and further network expansion.

Concurrent Learning. Stage 1 (Figure 1a,b) utilizes our MLP to drive Pynta's exhaustive TS exploration module. This exploration is leveraged to iteratively train an rMLP to learn the reaction coordinate region for a set of input reactions. As the details of Pynta⁴⁸ are discussed in prior work, only a cursory overview is provided here. Specifically, we first use our MLP to drive the harmonically forced saddle point (HFSP) method⁴⁸ to generate a comprehensive set of transition state guesses corresponding to each elementary reaction in our input mechanism. This step is crucial, as HFSP considers all unique site-pair combinations of the reactants or products to produce an exhaustive list of possible TS guesses. Note that the reactant and product geometries, a prerequisite for the HFSP algorithm, are sourced from Pynta's adsorbate optimization step,⁴⁸ where Pynta attempts to optimize each unique adsorbate on all unique sites. To obtain a tractable set of transition states, we apply Pynta's harmonic filtering to remove overly strained structures that are likely to be either invalid or high-energy saddle points. For example, for the oxygen-assisted methanol oxidation step (reaction 1 in Table S1 and the example shown in Figure 1c), HFSP begins with 129 TS guess geometries for the hydrogen atom transfer reaction from methanol to surface-bound oxygen. After harmonic filtering and similarity analysis, 34 structures are retained. Figure S1 shows the number of HFSP configurations considered for all 53 reactions (~1,100 filtered HFSP configurations across the reaction set). The importance of considering many possible transition state guesses is demonstrated in the 0.25 eV range of transition state energies determined (Figure 1c).

Each filtered transition state guess is then optimized using our MLP to run Sella, our efficient optimizer.^{42–44} Sella constructs an approximate Hessian with our MLP to optimize the guess structures to a valid TS, i.e., first-order saddle points on the underlying potential energy surface. If the MLP/TS is successfully obtained, then we utilize the intrinsic reaction coordinate (IRC) method to confirm that the MLP/TS geometry connects the intended reactants and products. For the purposes of this study, we augmented the Pynta workflow with the ability to run short, 0.5 ps NVT Langevin molecular dynamics simulations at 298 K using a 0.5 fs time step to further expand the diversity of our dataset. Although MD simulations are not required to obtain the transition states, previous work from us and others has shown that MD data can improve the quality of the resulting MLP model.⁶⁸ Furthermore, since MD is performed using the MLP, there is only an incremental increase in the required computational cost. Specifically, MLP/MD runs are conducted starting from harmonically filtered MLP/HFSP configurations in the initial iterations and from Sella-refined MLP/TS configurations in later iterations (see Figure S2).

The simulations described above comprise the exploration block of stage 1. At each iteration of our workflow, the explored configurations are downsampled using our ensemble-based uncertainty metric ϵ_t (eq 1). Sampled configurations are subsequently labeled with DFT single-point (DFT/SPE) calculations. We employ a sampling approach which uniformly selects configurations within a given uncertainty range ($\epsilon_{t,\min} < \epsilon_t < \epsilon_{t,\max}$ broken into 10 bins), as detailed in the Supporting Information, to select a diverse range of configurations that

encompass a spectrum of model uncertainty. The sampling is also dynamic so that more configurations are labeled per bin from reaction samplers (i.e., Sella, MD, IRC) which have a higher max uncertainty (see Tables S2 and S3).

Transition State Auditing. The concurrent learning algorithm outlined above generates a set of TS geometries for each elementary step through various iterations of the rMLP model. To ensure that these TS guesses correspond to the correct reactions, we developed a multistep auditing protocol (Figure 2). This TS auditing stage automatically

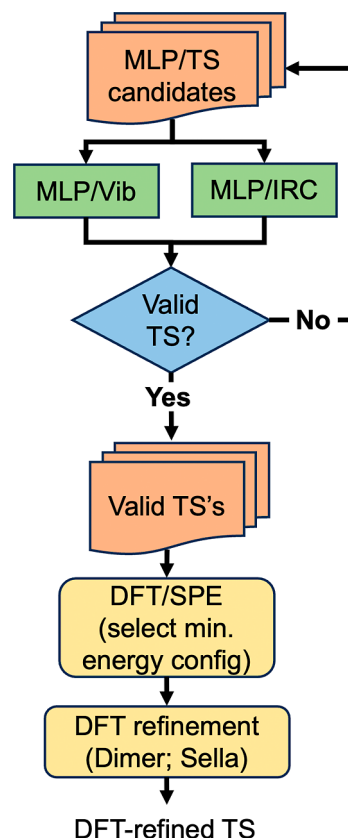


Figure 2. Overview of stage 2 and stage 3, the transition state auditing and refinement workflow. Each candidate MLP/TS for a given reaction is evaluated based upon MLP/vib and MLP/IRC. Following evaluation, the set of valid TSs is evaluated with DFT/SPE and sorted by DFT energy. The minimum energy valid TS moves on to stage 3 DFT refinement.

analyzes a set of MLP-optimized transition states based on several criteria to confirm that the most stable (i.e., lowest energy) and valid TS structure is used for further refinement in stage 3. In addition to verifying that each TS has a single imaginary frequency, we confirm that the MLP/TS connects the intended reactants and products. Specifically, the adjacency lists of the reactants and products supplied to Pynta are automatically compared against the end points of the IRC. This is done by converting these IRC end points into a two-dimensional representation and performing a graph isomorphism test with the Pynta input. This tool was built upon previous tools developed to study coverage dependence⁷⁴ within Pynta. An example of the TS auditing algorithm is provided in Supporting Information (Figure S3). We sort the final set of valid MLP/TSs by energy using a DFT single-point energy calculation. In this work, the minimum energy valid

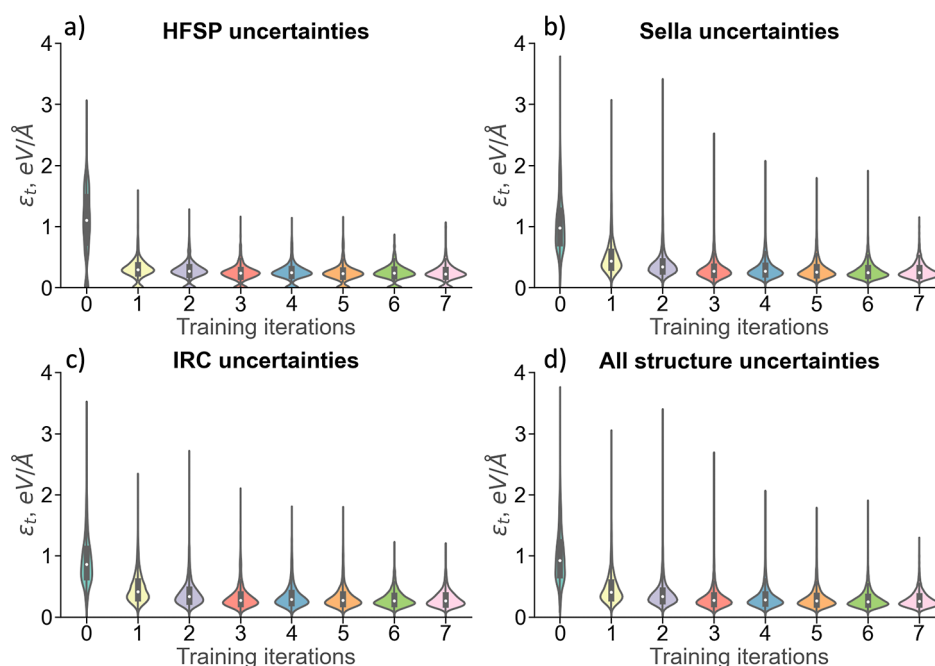


Figure 3. Model uncertainties ϵ_t for each model DP_i , $i = 1, \dots, 7$ on the entire data set generated over 7 training iterations in the concurrent learning loop.

MLP/TS is selected for further DFT refinement. We note that accurate selection of the minimum energy TS requires that the MLP predictions are sufficiently accurate. Depending on the model accuracy obtained through concurrent learning, additional transition state structures can be selected for refinement to confidently identify the minimum energy TS configuration. As will be discussed in subsequent sections, we completed 7 iterations of concurrent learning before moving to transition state auditing and refinement.

Transition State Refinement. Stage 3 is important to ensure that we maintain the DFT-level accuracy. As demonstrated below, this three-stage workflow effectively circumvents the most computationally expensive steps of the original Pynta/DFT workflow, delivering DFT-quality results faster. After auditing, stage 3 uses a relatively quick DFT/dimer⁷⁵ calculation to optimize the MLP/TS to a DFT saddle point. If DFT/dimer proves unsuccessful, we then use DFT/Sella. The resulting DFT/TSs are confirmed through automated DFT-vibrational analysis.

Model Performance. We implemented the above workflow to investigate the partial oxidation of methanol over a Ag(111) catalyst. We initialized our workflow using a DeepMD-kit model^{66,67,70} previously trained to study adsorbate diffusion on Ag(111).⁶² Beyond this iteration 0 model (denoted as DP_0), the inputs to the rMLP curriculum include (a) a predefined set of elementary reactions and (b) the corresponding DFT-optimized Ag(111) slab and adsorbate configurations generated by Pynta. Our reaction set is the 53 elementary reactions (Table S1) previously described by Aljama et al.,²⁸ which encompasses a diverse range of chemistry, including more than 20 unique surface-bound species, thus presenting a significant challenge to test the efficacy of our curriculum. To evaluate the accuracy and computational cost of our approach, we also calculated a set of 23 reaction barriers using Pynta/DFT only. The computational costs associated with a few DFT-only studies are summarized in Table S4.

The workflow in Figure 1 was employed in seven concurrent learning iterations to sample configurations along the reaction coordinate regions in parallel for a set of 53 reactions. The resulting iteration 7 model (denoted as DP_7) has been trained on a total of 26,483 DFT-labeled HFSP, Sella, IRC, and MD configurations. A detailed breakdown of the training data set is presented in Figure S2 and the training data is available as an ase database in the GitHub repository provided in the Supporting Information.

Figure 3a illustrates the model uncertainties across the entire DFT data set for all 53 reactions obtained over seven concurrent training iterations. In this analysis, we focus on configurations with a maximum force on each atom below 2 eV/Å, as these are more likely to yield valid transition state geometries. Notably, while the median values of ϵ_t drop rapidly to below 0.5 eV/Å after the very first iteration, the tails of the distributions at higher ϵ_t values decrease at a more gradual rate.

Specifically, we observe a significant reduction in our model's uncertainty for HFSP configurations (Figure 3b) compared to those for Sella (Figure 3c), IRC (Figure 3d), and MD (Figure S4). This reduction can be attributed to two main factors: (1) all HFSP configurations until iteration 2 are labeled using DFT and (2) the harmonic potentials used to generate the HFSP configurations result in configurations that are relatively consistent throughout the training process. Utilizing all HFSP configurations for DFT labeling in the early model iterations is particularly advantageous, as Sella, our TS optimizer, relies on finite differences to generate approximate Hessians to optimize HFSP TS guesses to a first-order saddle point.

In contrast to the favorable trends observed for HFSP configurations, those derived from Sella and IRC configurations exhibit greater diversity and show slow yet consistent improvements over the seven training iterations. Specifically, several additional iterations are needed to reduce the outliers at higher ϵ_t values to below 1.0 eV/Å for the Sella trajectories. A detailed investigation into this issue reveals that some

configurations remain challenging to describe even with the DP₇ model. These problematic configurations often arise from unphysical HFSP guesses, which can persist during the training protocol. Recently, overcoming this issue has been explored using improved HFSP algorithms to obtain more accurate TS guesses.⁷⁶

After seven training iterations, the DP₇ model generated valid MLP/TSs for 51 of the 53 reactions. While transition states for two reactions (i.e., reactions 19 and 53) were not obtained using DP₇, surprisingly, we found that previous model iterations were successful for these two cases. Furthermore, four reactions (Rxn 8, 12, 19, and 27) failed at the DFT-refinement step, representing a 10% failure rate across the entire workflow. The underlying cause for these failed reactions is discussed later.

In Figure 4, we compare the energetics of our rMLP-calculated barriers to our validation set of 23 Pynta-DFT

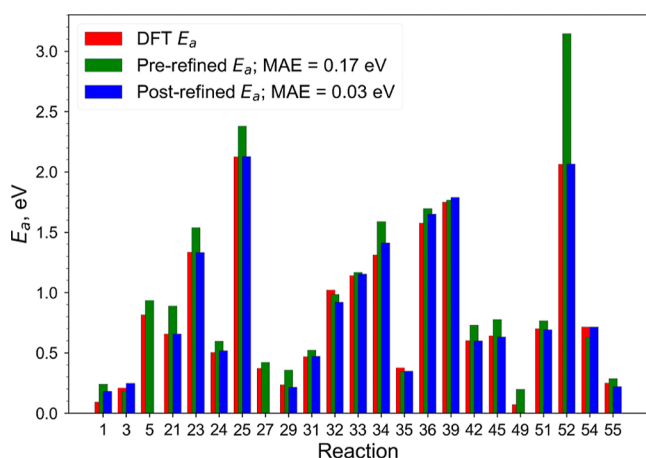


Figure 4. Comparison of activation energies, E_a , from pure DFT-Pynta (red), pre-refined minimum energy MLP/TS + DFT/SPE (green), and post-refined MLP/TS + DFT/refinement (blue).

barriers. We generally observe that the pre-refined energy barriers (green) are slightly higher than the Pynta/DFT (red) ground truth. Indeed, the mean-absolute errors (MAE) reduce from 0.17 eV (prerefinement) to 0.03 eV after stage 3 refinement across the 23 validation reactions.

Before proceeding further, it is useful to discuss the computational cost associated with this approach. Overall we obtain a 7× speedup compared to base Pynta for the set of 53 reactions. A complete accounting of the computational expense is provided in the Supporting Information (Table S5). One key metric for model performance is the inference time, which is the amount of time needed for the MLP to predict the energy and forces on a new atomic configuration. A benefit of our model is that fast inference times (on the order of milliseconds on a single CPU node) enable the Pynta workflow to be completed in just a few hours and can allow the model to be used for subsequent dynamic studies. Also, despite the relatively large size of our DFT data set, we are able to train our model in less than 8 h on a single CPU node on the NERSC Perlmutter machine (2× AMD EPYC 7763). Achieving DFT accuracy for transition states with this level of speedup facilitates larger-scale studies than were previously feasible due to computational constraints. We note that even greater speedups are possible, and we will discuss potential improvements to this workflow in subsequent sections.

Curriculum Limitations and Potential Improvements.

At this point, we turn our attention back to the model shortcomings identified previously. Specifically, we focus on understanding (1) why earlier model iterations are sometimes better at finding certain MLP/TSs and (2) why does the DFT-refinement step fail for certain reactions?

To understand the model degradation observed for reactions 19 and 53, we analyzed the trends in model performance for 53 reactions across the seven concurring learning iterations. We observed that although the total model uncertainty consistently improves (Figure 3) during the training protocol, the model performance for some reactions shows interesting trends. For instance, O₂ dissociation, as shown in Figure 5, exhibits decreasing uncertainties until iteration three, after which the

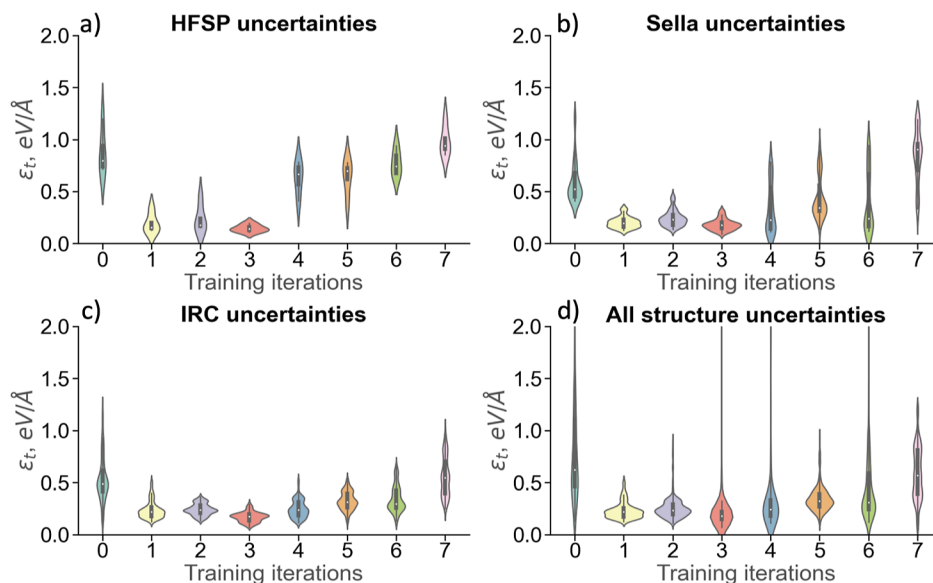


Figure 5. Model uncertainties ϵ_t for each model DP_{*i*}, *i* = 1,..., 7 on the configurations sampled for reaction 53, O₂ + * + * ⇌ O* + O*, over seven training iterations in the concurrent learning loop.

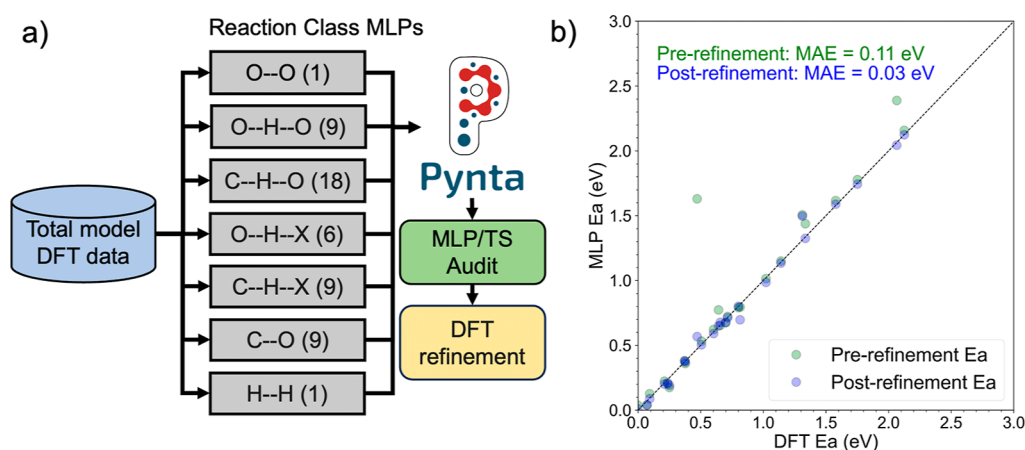


Figure 6. Description of reaction class categories and associated number of reactions in panel (a). Parity plot of MLP and DFT reaction barriers in panel (b).

uncertainty worsens across all samplers (Figure 5a–d). As a consequence, the DP₃ model successfully identified an MLP/TS that was successfully refined using DFT. In contrast, the DP₇ model did not yield a valid MLP/TS for O₂ dissociation.

This observation underscores a key limitation of the parallel training approach when it is combined with an adaptive sampling scheme. Specifically, while the DP₃ model successfully captures the relevant portion of the reaction coordinate for O₂ dissociation, the adaptive sampling scheme induces an on-the-fly data bias as configurations relevant to this reaction are sampled less as the uncertainty decreased (see Figure S6). This issue is further exacerbated for the dissociation of the O₂ bond, as the reaction network does not contain any other O–O bond formation reactions. While not discussed here, similar trends are observed for reaction 19 (Figure S7).

Building on the above understanding regarding stage 2 failures, we now investigate four reactions that failed during the DFT refinement in stage 3. Hypothesizing that suboptimal model fitting of the DFT data could be a possible cause, we first used hyperparameter optimization to improve the model performance. Specifically, in addition to exploring different sizes of the fitting and embedding nets,⁶⁸ we created two different models with varying values of key hyperparameters (Tables S7 and S8).⁷⁷ Although this approach reduced the mean squared error for forces on the training data from 0.018 to 0.013 eV/Å, the hyperparameter-optimized model still failed during the DFT-refinement stage for all four reactions. Thus, this suggests that similar to the above total model analysis, data bias could also be a contributing factor to these failures.

To circumvent this challenge, we explored a reaction-class-based training strategy that categorized the 53 reactions into multiple distinct groups based on the types of bonds being broken. This categorization resulted in seven different models (N.B.: not to be confused with the seven previous model iterations), each corresponding to a specific category of transition state (Figure 6a). We note here that there are many possible rules that could have been applied to demarcate different reaction classes. Exploration of the optimal classification scheme is outside the scope of this current work. However, the underlying philosophy of information being shared across different reactions (within a given class) will likely be helpful for developing accurate MLPs on diverse reaction networks. We believe that our classification scheme is simple and generalizable to other catalytic systems and

additional heuristics can be applied for potential edge cases where the rules do not uniquely classify a reaction. Each model was trained with the base data set from the DP₀ model, along with all data generated from Pynta for each reaction class over the 7 concurrent learning iterations from the total model approach. Compared with the total model approach discussed above, the reaction-class strategy successfully identified transition states (TS) for *all* 53 reactions, both before and after DFT refinement. As shown in Figure 6b, the mean absolute error (MAE) for pre-refinement TSs was 0.11 eV, which improved to 0.03 eV after post-refinement. Additionally, the reaction class models exhibited lower mean absolute errors on force predictions when compared with the total model, even after hyperparameter tuning of the latter. The vibrational frequencies of the minimum energy transition states are provided for reference in Table S9. We gain additional insights when investigating the reactions that failed DFT refinement in our total approach but were successful in the reaction class approach. As shown in Figure S8, the reaction class approach obtained lower model uncertainties and DFT-residual forces compared to the total model approach in most cases.

As this reaction class approach provides a comprehensive exploration of all transition state geometries, we now compare the results of this automated workflow with those of previous studies in the literature. Specifically, previous work by Aljama et al.²⁸ has reported the free energy diagram of the minimum energy pathway (MEP) connecting methanol (i.e., the reactant) to the desired product (i.e., formaldehyde). Figure 7 compares these literature results to our MEP obtained using the DFT-refined transition states from our reaction class model. In general, we observe that the energy barriers obtained using the Pynta curriculum are as much as 0.4 eV lower than those from Aljama et al.²⁸ While differences in the DFT functional of choice (i.e., RPBE-D3 vs BEEF-vdW)^{78,79} between the two studies could explain some of the differences, these results demonstrate the importance of obtaining the lowest-energy transition states. We note that comparing our optimized geometries with those of Aljama et al. was not possible, since the transition state geometries for their calculations were not reported.

Generalization to Large Pretrained Models. As an alternative acceleration approach to our concurrent learning strategy for training a system-specific model, we explore the use of a large pretrained OC20³⁹ model. We focus on the

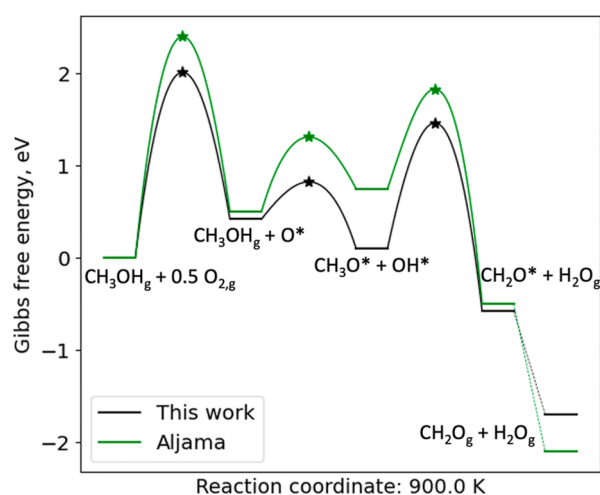


Figure 7. Minimum free energy diagram of formaldehyde synthesis from methanol partial oxidation on Ag(111).

EquiformerV2⁴⁰ model, which has demonstrated exceptional performance on the OC20 data set and has recently excelled in transition state optimization⁴¹ and Hessian calculations.⁴⁵ Initial tests with the pretrained EquiformerV2 models (eq 31M and eq 153M parameters) indicated a need for fine-tuning due to differences in DFT parameters used for optimizing our adsorbates and slab compared to the ones to which the original model was fit, particularly the inclusion of van der Waals corrections^{59,61} (see Figures S9–S12). The motivation behind the inclusion of such long-range corrections comes from many works demonstrating poor description of physisorbed species with RPBE^{78,80} and the enhanced description with D3 corrections.⁷⁹ We fine-tuned the eq 31M and eq 153M models on 15,234 DFT/SPEs generated in the study of adsorbate diffusion on Ag(111) from our base DP₀ model.⁶² Details of our train–test–validation split are provided in Supporting Information. Given the similar performance and faster inference time of the eq 31M model on our test reaction, we ran Pynta with the eq 31M model on our 53 reactions. After MLP auditing, we identified 47 valid MLP/TSs. The maximum residual forces (Figure S13) for the OCP model align closely with those from the reaction class approach, suggesting that these transition states would succeed during DFT refinement. However, since we have already obtained these transition states, further refinement was not conducted. Figure 8 shows that when compared to the DFT-calculated barriers, the tuned eq 31M model achieves great parity with DFT. When investigating the five reactions for which we failed to find an MLP/TS with the eq 31M model, we found that many Sella saddle-point optimizations failed to converge due to an inability of the model to continuously reduce the forces. Recent work has suggested that fine-tuning of the OCP model on finite-difference data can improve the model's accuracy in calculating Hessians. The OC20³⁹ data set only included optimizations where step sizes are much larger than what is required to calculate Hessians, so the inclusion of data points with small displacements could help the model capture the curvature information on the PES. While further DFT refinement and fine-tuning of the OCP approach is outside the scope of this study, these encouraging results motivate the advantages of combining Pynta with pretrained models including those based on equivariant architectures.

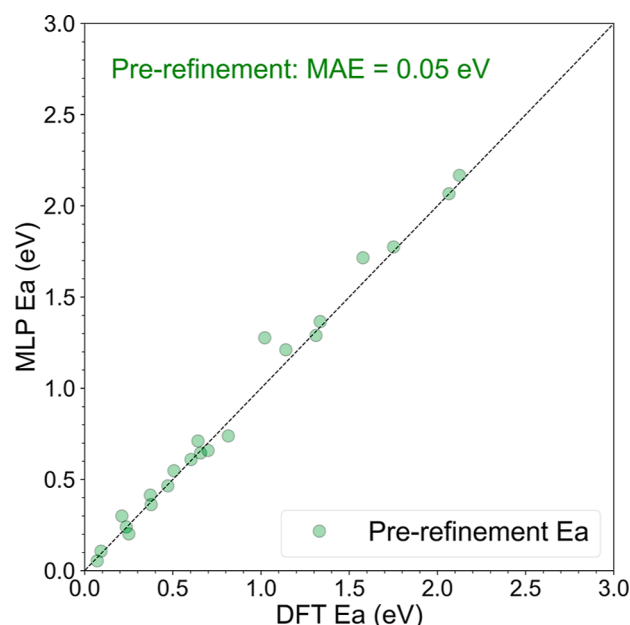


Figure 8. Parity plot of MLP and DFT-reaction barriers using a fine-tuned OC20-pretrained EquiformerV2 31 M model.

Comparison of Various Approaches. Table 1 shows the results of the various model approaches used in this study. The

Table 1. Results of DFT TS Refinement for Each Approach

approach	number of valid MLP-TS	number of valid DFT-refined TS	computational expense (total ^a)
total model (1)	53	49	3503
reaction class models (7)	53	53	3569
EquiformerV2-31M ^b	47	-	1092
DFT-only Pynta ^c	-	-	25,007

^aTotal time (CPU hours) utilized on a single CPU node on a Perlmutter machine (2× AMD EPYC 7763). ^bInitially trained for S2EF-OC20-All + MD. ^cThe computational cost for a few Pynta-DFT reactions is shown in Table S4.

reaction class model achieved a 7× speedup compared to the base Pynta/DFT approach and was successful for all reactions pre- and post-refinement. We did include a computational expense for the eq 31M model, although we did not proceed with DFT refinement of these transition states. Without DFT refinement, the eq 31M model achieved a 20× speedup with an 89% success rate. Both approaches have advantages and disadvantages as discussed previously. While large pretrained models can provide greater acceleration for optimization tasks, the inference time for these models is several orders of magnitude slower than smaller, system-specific models. There is a trade-off in initial acceleration and extensibility to explore system dynamics, and future work should investigate options for leveraging both model types. Overall our results demonstrate the considerable speedup in Pynta's comprehensive TS exploration that can be achieved using our concurrent learning loop or fine-tuning of a pretrained MLP.

CONCLUSIONS

This work demonstrates how Pynta, an automated kinetics workflow tool, can be combined with reactive machine learning

potentials (rMLPs) to accelerate the calculation of reaction barriers and transition state geometries for several surface-mediated reactions, simultaneously. Specifically, by using the industrially relevant process of methanol partial oxidation catalyzed by Ag(111) as an example, we showcase an automated concurrent learning protocol to train a system-specific rMLP on a diverse set of 53 reactions. The total model obtained a 92% success rate, but investigations into the failed reactions revealed that data bias may have resulted in model degradation. Next, to overcome these limitations, we used a reaction class approach that outperformed the total model approach, and we obtained valid transition states from the entire set of 53 reactions. By comparing these results with a pure Pynta/DFT strategy, we show that our reaction class approach is seven times faster and finds essentially identical activation barriers. We note that the accuracy of the overall kinetics predicted by using this workflow depends on the reliability of the activation energies calculated for the rate-determining step. Careful analysis of specific reactions is required during model training. Depending on the model accuracy achieved, additional transition states should be considered for DFT refinement to ensure that the minimum energy transition state is obtained for the kinetically relevant reactions. Finally, we show that Pynta can be accelerated with a fine-tuned foundational model. Specifically, we demonstrate that fine-tuning the OC20 pretrained EquiformerV2 (eq 31M) model results in an 89% success rate and achieves excellent parity with our validation set of reaction barriers (MAE = 0.05 eV) without further transition state refinement. The pretrained model also achieves an impressive 20× speedup. All of the tools developed in this work are now available for use with Pynta, which can aid in accelerating the exploration of minimum energy pathways relevant to generating kinetic models for surface-mediated reactions. Automatic generation of thermochemistry and kinetic parameters for use in the RMG can enable further reaction expansion. Current work in the group is focused on exploring strategies to overcome data bias issues, generalizing Pynta-based curricula to other rMLP architectures, and expanding these tools to describe other catalysts.

■ ASSOCIATED CONTENT

Data Availability Statement

The following are available on GitHub: training data, DP₇ model, reaction class models, database of DFT-refined transition state configurations <https://github.com/tdprice-858/Pynta-ML>.

SI Supporting Information

The Supporting Information is available free of charge at <https://pubs.acs.org/doi/10.1021/acs.jpcc.5c00305>.

Reaction list, supplementary results, and computational costs (PDF)

■ AUTHOR INFORMATION

Corresponding Authors

Judit Zádor – Combustion Research Facility, Sandia National Laboratories, Livermore, California 94550, United States;
orcid.org/0000-0002-9123-8238; Email: jzador@sandia.gov

Ambarish Kulkarni – Department of Chemical Engineering, University of California, Davis, California 95616, United

States; orcid.org/0000-0001-9834-8264;

Email: arkulkarni@ucdavis.edu

Authors

Trevor Price – Department of Chemical Engineering, University of California, Davis, California 95616, United States

Saurabh Sivakumar – Department of Chemical Engineering, University of California, Davis, California 95616, United States

Matthew S. Johnson – Combustion Research Facility, Sandia National Laboratories, Livermore, California 94550, United States

Complete contact information is available at:

<https://pubs.acs.org/10.1021/acs.jpcc.5c00305>

Notes

The authors declare no competing financial interest.

■ ACKNOWLEDGMENTS

TDP was partially supported through a DOE Office of Science Graduate Student Research (SCGS) Program award. AK, SS, and TDP were supported by the U.S. Department of Energy (DOE), Office of Basic Energy Sciences (BES), Division of Chemical Sciences, Geosciences and Biosciences (CSGB) under Grant DE-SC0020320. MSJ and JZ were supported by the Exascale Catalytic Chemistry (ECC) Project, which is funded by the U.S. Department of Energy, Office of Science, Basic Energy Sciences, Chemical Sciences, Geosciences and Biosciences Division, as part of the Computational Chemistry Sciences Program. This research used resources of the National Energy Research Scientific Computing Center (NERSC), a Department of Energy Office of Science User Facility using NERSC award BES-ERCAP0026789 and NERSC award BES-ERCAP0028388. This article has been authored by employees of National Technology & Engineering Solutions of Sandia, LLC under Contract No. DE-NA0003525 with the U.S. Department of Energy (DOE). The employees co-own right, title, and interest in and to the article and are responsible for its contents. The United States Government retains and the publisher, by accepting the article for publication, acknowledges that the United States Government retains a nonexclusive, paid-up, irrevocable, worldwide license to publish or reproduce the published form of this article or allow others to do so, for United States Government purposes. DOE will provide public access to these results of federally sponsored research in accordance with the DOE Public Access Plan (<https://www.energy.gov/downloads/doe-public-access-plan>).

■ REFERENCES

- (1) Motagamwala, A. H.; Dumesic, J. A. Microkinetic Modeling: A Tool for Rational Catalyst Design. *Chem. Rev.* **2021**, *121*, 1049–1076.
- (2) Chen, B. W. J.; Xu, L.; Mavrikakis, M. Computational Methods in Heterogeneous Catalysis. *Chem. Rev.* **2021**, *121*, 1007–1048.
- (3) Kreitz, B.; Lott, P.; Bae, J.; Blöndal, K.; Angeli, S.; Ulissi, Z. W.; Studt, F.; Goldsmith, C. F.; Deutschmann, O. Detailed Microkinetics for the Oxidation of Exhaust Gas Emissions through Automated Mechanism Generation. *ACS Catal.* **2022**, *12*, 11137–11151.
- (4) Ulissi, Z. W.; Medford, A. J.; Bligaard, T.; Nørskov, J. K. To address surface reaction network complexity using scaling relations machine learning and DFT calculations. *Nat. Commun.* **2017**, *8*, 14621.

- (5) Shambhawi, Mohan, O.; Choksi, T. S.; Lapkin, A. A. The design and optimization of heterogeneous catalysts using computational methods. *Catal. Sci. Technol.* **2024**, *14*, 515–532.
- (6) Neugeboren, J.; Borodin, D.; Hahn, H. W.; Altschäff, J.; Kandratenka, A.; Auerbach, D. J.; Campbell, C. T.; Schwarzer, D.; Harding, D. J.; Wodtke, A. M.; et al. Velocity-resolved kinetics of site-specific carbon monoxide oxidation on platinum surfaces. *Nature* **2018**, *558*, 280–283.
- (7) Zhou, B.; Huang, E.; Almeida, R.; Gurses, S.; Ungar, A.; Zetterberg, J.; Kulkarni, A.; Kronawitter, C. X.; Osborn, D. L.; Hansen, N.; et al. Near-Surface Imaging of the Multicomponent Gas Phase above a Silver Catalyst during Partial Oxidation of Methanol. *ACS Catal.* **2021**, *11*, 155–168.
- (8) Pfaff, S.; Huang, E.; Frank, J. H. Imaging Gas-Phase Methyl Radicals over a Ag/SiO₂ Catalyst during the Partial Oxidation of Methanol. *ACS Catal.* **2024**, *14*, 15590–15598.
- (9) Zetterberg, J.; Blomberg, S.; Gustafson, J.; Sun, Z. W.; Li, Z. S.; Lundgren, E.; Aldén, M. An in situ set up for the detection of CO₂ from catalytic CO oxidation by using planar laser-induced fluorescence. *Rev. Sci. Instrum.* **2012**, *83*, 053104.
- (10) Matera, S.; Blomberg, S.; Hoffmann, M. J.; Zetterberg, J.; Gustafson, J.; Lundgren, E.; Reuter, K. Evidence for the Active Phase of Heterogeneous Catalysts through In Situ Reaction Product Imaging and Multiscale Modeling. *ACS Catal.* **2015**, *5*, 4514–4518.
- (11) Blomberg, S.; Brackmann, C.; Gustafson, J.; Aldén, M.; Lundgren, E.; Zetterberg, J. Real-Time Gas-Phase Imaging over a Pd(110) Catalyst during CO Oxidation by Means of Planar Laser-Induced Fluorescence. *ACS Catal.* **2015**, *5*, 2028–2034.
- (12) Blomberg, S.; Zhou, J.; Gustafson, J.; Zetterberg, J.; Lundgren, E. 2D and 3D imaging of the gas phase close to an operating model catalyst by planar laser induced fluorescence. *J. Phys.: Condens. Matter* **2016**, *28*, 453002.
- (13) Blomberg, S.; Zetterberg, J.; Zhou, J.; Merte, L. R.; Gustafson, J.; Shipilin, M.; Trinchero, A.; Miccio, L. A.; Magaña, A.; Ilyn, M.; et al. Strain Dependent Light-off Temperature in Catalysis Revealed by Planar Laser-Induced Fluorescence. *ACS Catal.* **2017**, *7*, 110–114.
- (14) Zhou, J.; Blomberg, S.; Gustafson, J.; Lundgren, E.; Zetterberg, J. Visualization of Gas Distribution in a Model AP-XPS Reactor by PLIF: CO Oxidation over a Pd(100) Catalyst. *Catalysts* **2017**, *7*, 29.
- (15) Zhou, J.; Matera, S.; Pfaff, S.; Blomberg, S.; Lundgren, E.; Zetterberg, J. Combining Planar Laser-Induced Fluorescence with Stagnation Point Flows for Small Single-Crystal Model Catalysts: CO Oxidation on a Pd(100). *Catalysts* **2019**, *9*, 484.
- (16) Zellner, A.; Suntz, R.; Deutschmann, O. Two-Dimensional Spatial Resolution of Concentration Profiles in Catalytic Reactors by Planar Laser-Induced Fluorescence: NO Reduction over Diesel Oxidation Catalysts. *Angew. Chem., Int. Ed.* **2015**, *54*, 2653–2655.
- (17) Wan, S.; Torkashvand, B.; Häber, T.; Suntz, R.; Deutschmann, O. Investigation of HCHO Catalytic Oxidation over Platinum using Planar Laser-Induced Fluorescence. *Appl. Catal., B* **2020**, *264*, 118473.
- (18) Xie, W.; Xu, J.; Chen, J.; Wang, H.; Hu, P. Achieving Theory–Experiment Parity for Activity and Selectivity in Heterogeneous Catalysis Using Microkinetic Modeling. *Acc. Chem. Res.* **2022**, *55*, 1237–1248.
- (19) Goldsmith, C. F.; West, R. H. Automatic Generation of Microkinetic Mechanisms for Heterogeneous Catalysis. *J. Phys. Chem. C* **2017**, *121*, 9970–9981.
- (20) Johnson, M. S.; Dong, X.; Grinberg Dana, A.; Chung, Y.; Farina, D., Jr.; Gillis, R.; Liu, M.; Yee, N.; Blondal, K.; Mazeau, E.; et al. RMG Database for Chemical Property Prediction. *J. Chem. Inf. Model.* **2022**, *62*, 4906–4915.
- (21) Mazeau, E. J.; Satpute, P.; Blöndal, K.; Goldsmith, C. F.; West, R. H. Automated Mechanism Generation Using Linear Scaling Relationships and Sensitivity Analyses Applied to Catalytic Partial Oxidation of Methane. *ACS Catal.* **2021**, *11*, 7114–7125.
- (22) Liu, M.; Grinberg Dana, A.; Johnson, M. S.; Goldman, M. J.; Jocher, A.; Payne, A. M.; Grambow, C. A.; Han, K.; Yee, N. W.; Mazeau, E. J.; et al. Reaction Mechanism Generator v3.0: Advances in Automatic Mechanism Generation. *J. Chem. Inf. Model.* **2021**, *61*, 2686–2696.
- (23) Gao, C. W.; Allen, J. W.; Green, W. H.; West, R. H. Reaction Mechanism Generator: Automatic construction of chemical kinetic mechanisms. *Comput. Phys. Commun.* **2016**, *203*, 212–225.
- (24) Johnson, M. S.; Pang, H.-W.; Liu, M.; Green, W. H. Species selection for automatic chemical kinetic mechanism generation. *Int. J. Chem. Kinet.* **2025**, *57*, 93–107.
- (25) Kreitz, B.; Wehinger, G. D.; Goldsmith, C. F.; Turek, T. Microkinetic Modeling of the Transient CO₂ Methanation with DFT-Based Uncertainties in a Berty Reactor. *ChemCatChem* **2022**, *14*, No. e202200570.
- (26) Blondal, K.; Jelic, J.; Mazeau, E.; Studt, F.; West, R. H.; Goldsmith, C. F. Computer-Generated Kinetics for Coupled Heterogeneous/Homogeneous Systems: A Case Study in Catalytic Combustion of Methane on Platinum. *Ind. Eng. Chem. Res.* **2019**, *58*, 17682–17691.
- (27) Evans, M. G.; Polanyi, M. Inertia and driving force of chemical reactions. *Trans. Faraday Soc.* **1938**, *34*, 11–24.
- (28) Aljama, H.; Yoo, J. S.; Nørskov, J. K.; Abild-Pedersen, F.; Studt, F. Methanol Partial Oxidation on Ag(111) from First Principles. *ChemCatChem* **2016**, *8*, 3621–3625.
- (29) Henkelman, G.; Uberuaga, B. P.; Jónsson, H. A climbing image nudged elastic band method for finding saddle points and minimum energy paths. *J. Chem. Phys.* **2000**, *113*, 9901–9904.
- (30) Henkelman, G.; Jónsson, H. Improved tangent estimate in the nudged elastic band method for finding minimum energy paths and saddle points. *J. Chem. Phys.* **2000**, *113*, 9978–9985.
- (31) Omranpour, A.; Elsner, J.; Lausch, K. N.; Behler, J. Machine Learning Potentials for Heterogeneous Catalysis. *arXiv* **2024**, arXiv:2411.00720.
- (32) Behler, J.; Parrinello, M. Generalized Neural-Network Representation of High-Dimensional Potential-Energy Surfaces. *Phys. Rev. Lett.* **2007**, *98*, 146401.
- (33) Peterson, A. A. Acceleration of saddle-point searches with machine learning. *J. Chem. Phys.* **2016**, *145*, 074106.
- (34) Yang, Y.; Jiménez-Negrón, O. A.; Kitchin, J. R. Machine-learning accelerated geometry optimization in molecular simulation. *J. Chem. Phys.* **2021**, *154*, 234704.
- (35) Garrido Torres, J. A.; Jennings, P. C.; Hansen, M. H.; Boes, J. R.; Bligaard, T. Low-Scaling Algorithm for Nudged Elastic Band Calculations Using a Surrogate Machine Learning Model. *Phys. Rev. Lett.* **2019**, *122*, 156001.
- (36) Schaaf, L. L.; Fako, E.; De, S.; Schäfer, A.; Csányi, G. Accurate energy barriers for catalytic reaction pathways: an automatic training protocol for machine learning force fields. *npj Comput. Mater.* **2023**, *9*, 180.
- (37) Klawohn, S.; Darby, J. P.; Kermode, J. R.; Csányi, G.; Caro, M. A.; Bartók, A. P. Gaussian approximation potentials: Theory, software implementation and application examples. *J. Chem. Phys.* **2023**, *159*, 174108.
- (38) Zuo, Y.; Chen, C.; Li, X.; Deng, Z.; Chen, Y.; Behler, J.; Csányi, G.; Shapeev, A. V.; Thompson, A. P.; Wood, M. A.; Ong, S. P. Performance and Cost Assessment of Machine Learning Interatomic Potentials. *J. Phys. Chem. A* **2020**, *124*, 731–745.
- (39) Chanussot, L.; Das, A.; Goyal, S.; Lavril, T.; Shuaibi, M.; Riviere, M.; Tran, K.; Heras-Domingo, J.; Ho, C.; Hu, W.; et al. Open Catalyst 2021 (OC20) Dataset and Community Challenges. *ACS Catal.* **2021**, *11*, 6059–6072.
- (40) Liao, Y.-L.; Wood, B. M.; Das, A.; Smidt, T. EquiformerV2: Improved Equivariant Transformer for Scaling to Higher-Degree Representations. *The Twelfth International Conference on Learning Representations*, 2024.
- (41) Wander, B.; Shuaibi, M.; Kitchin, J. R.; Ulissi, Z. W.; Zitnick, C. L. CatTSunami: Accelerating Transition State Energy Calculations with Pre-trained Graph Neural Networks. *ACS Catal.* **2025**, *15*, 5283–5294.

- (42) Hermes, E. D.; Sargsyan, K.; Najm, H. N.; Zádor, J. Accelerated Saddle Point Refinement through Full Exploitation of Partial Hessian Diagonalization. *J. Chem. Theory Comput.* **2019**, *15*, 6536–6549.
- (43) Hermes, E. D.; Sargsyan, K.; Najm, H. N.; Zádor, J. Geometry optimization speedup through a geodesic approach to internal coordinates. *J. Chem. Phys.* **2021**, *155*, 094105.
- (44) Hermes, E. D.; Sargsyan, K.; Najm, H. N.; Zádor, J. Sella, an Open-Source Automation-Friendly Molecular Saddle Point Optimizer. *J. Chem. Theory Comput.* **2022**, *18*, 6974–6988.
- (45) Wander, B.; Musielewicz, J.; Cheula, R.; Kitchin, J. R. Accessing Numerical Energy Hessians with Graph Neural Network Potentials and Their Application in Heterogeneous Catalysis. *J. Phys. Chem. C* **2025**, *129*, 3510–3521.
- (46) Musielewicz, J.; Wang, X.; Tian, T.; Ulissi, Z. FINETUNA: fine-tuning accelerated molecular simulations. *Mach. Learn. Sci. Technol.* **2022**, *3*, 03LT01.
- (47) Wang, X.; Musielewicz, J.; Tran, R.; Kumar Ethirajan, S.; Fu, X.; Mera, H.; Kitchin, J. R.; Kurchin, R. C.; Ulissi, Z. W. Generalization of graph-based active learning relaxation strategies across materials. *Mach. Learn.: Sci. Technol.* **2024**, *5*, 025018.
- (48) Johnson, M. S.; Gierada, M.; Hermes, E. D.; Bross, D. H.; Sargsyan, K.; Najm, H. N.; Zádor, J. Pynta - An Automated Workflow for Calculation of Surface and Gas–Surface Kinetics. *J. Chem. Inf. Model.* **2023**, *63*, 5153–5168.
- (49) Lakshminarayanan, B.; Pritzel, A.; Blundell, C. Simple and Scalable Predictive Uncertainty Estimation using Deep Ensembles. *Advances in Neural Information Processing Systems*, 2017.
- (50) Yang, M.; You, R.; Li, D.; Zhang, Z.; Huang, W. Methanol Partial Oxidation Over Shaped Silver Nanoparticles Derived from Cubic and Octahedral Ag₂O Nanocrystals. *Catal. Lett.* **2019**, *149*, 2482–2491.
- (51) Millar, G. J.; Collins, M. Industrial Production of Formaldehyde Using Polycrystalline Silver Catalyst. *Ind. Eng. Chem. Res.* **2017**, *56*, 9247–9265.
- (52) Lervold, S.; Arnesen, K.; Beck, N.; Lødeng, R.; Yang, J.; Bingen, K.; Skjelstad, J.; Vennik, H. J. Morphology and Activity of Electrolytic Silver Catalyst for Partial Oxidation of Methanol to Formaldehyde Under Different Exposures and Oxidation Reactions. *Top. Catal.* **2019**, *62*, 699–711.
- (53) Sauer, J. The future of computational catalysis. *J. Catal.* **2024**, *433*, 115482.
- (54) Sauer, J. Ab Initio Calculations for Molecule–Surface Interactions with Chemical Accuracy. *Acc. Chem. Res.* **2019**, *52*, 3502–3510.
- (55) Hafner, J.; Kresse, G. *Properties of Complex Inorganic Solids*; Gonis, A., Meike, A., Turchi, P. E. A., Eds.; Springer US: Boston, MA, 1997; pp 69–82.
- (56) Kresse, G.; Joubert, D. From ultrasoft pseudopotentials to the projector augmented-wave method. *Phys. Rev. B* **1999**, *59*, 1758–1775.
- (57) Larsen, A. H.; Mortensen, J. J.; Blomqvist, J.; Castelli, I. E.; Christensen, R.; Dulak, M.; Friis, J.; Groves, M. N.; Hammer, B.; Hargus, C.; et al. The atomic simulation environment—a Python library for working with atoms. *J. Phys.: Condens. Matter* **2017**, *29*, 273002.
- (58) Hammer, B.; Hansen, L. B.; Nørskov, J. K. Improved adsorption energetics within density-functional theory using revised Perdew-Burke-Ernzerhof functionals. *Phys. Rev. B* **1999**, *59*, 7413–7421.
- (59) Smith, D. G. A.; Burns, L. A.; Patkowski, K.; Sherrill, C. D. Revised Damping Parameters for the D3 Dispersion Correction to Density Functional Theory. *J. Phys. Chem. Lett.* **2016**, *7*, 2197–2203.
- (60) Grimme, S.; Antony, J.; Ehrlich, S.; Krieg, H. A consistent and accurate ab initio parametrization of density functional dispersion correction (DFT-D) for the 94 elements H–Pu. *J. Chem. Phys.* **2010**, *132*, 154104.
- (61) Grimme, S.; Ehrlich, S.; Goerigk, L. Effect of the damping function in dispersion corrected density functional theory. *J. Comput. Chem.* **2011**, *32*, 1456–1465.
- (62) Sivakumar, S.; Kulkarni, A. Toward an ab Initio Description of Adsorbate Surface Dynamics. *J. Phys. Chem. C* **2024**, *128*, 13238–13248.
- (63) Davey, W. P. Precision Measurements of the Lattice Constants of Twelve Common Metals. *Phys. Rev.* **1925**, *25*, 753–761.
- (64) Monkhorst, H. J.; Pack, J. D. Special points for Brillouin-zone integrations. *Phys. Rev. B* **1976**, *13*, 5188–5192.
- (65) Zhang, L.; Han, J.; Wang, H.; Car, R.; E, W. Deep Potential Molecular Dynamics: A Scalable Model with the Accuracy of Quantum Mechanics. *Phys. Rev. Lett.* **2018**, *120*, 143001.
- (66) Zhang, L.; Han, J.; Wang, H.; Saidi, W.; Car, R.; Weinan, E. End-to-end Symmetry Preserving Inter-atomic Potential Energy Model for Finite and Extended Systems. *NeurIPS Proceedings; Advances in Neural Information Processing Systems*, 2018.
- (67) Zeng, J.; Zhang, D.; Lu, D.; Mo, P.; Li, Z.; Chen, Y.; Rynik, M.; Huang, L.; Li, Z.; Shi, S.; et al. DeePMD-kit v2: A software package for deep potential models. *J. Chem. Phys.* **2023**, *159*, 054801.
- (68) Guo, J.; Sours, T.; Holton, S.; Sun, C.; Kulkarni, A. R. Screening Cu-Zeolites for Methane Activation Using Curriculum-Based Training. *ACS Catal.* **2024**, *14*, 1232–1242.
- (69) Sours, T. G.; Kulkarni, A. R. Predicting Structural Properties of Pure Silica Zeolites Using Deep Neural Network Potentials. *J. Phys. Chem. C* **2023**, *127*, 1455–1463.
- (70) Zhang, Y.; Wang, H.; Chen, W.; Zeng, J.; Zhang, L.; Wang, H.; E, W. DP-GEN: A concurrent learning platform for the generation of reliable deep learning based potential energy models. *Comput. Phys. Commun.* **2020**, *253*, 107206.
- (71) Ruscic, B.; Pinzon, R. E.; Morton, M. L.; von Laszewski, G.; Bittner, S. J.; Nijssure, S. G.; Amin, K. A.; Minkoff, M.; Wagner, A. F. Introduction to Active Thermochemical Tables: Several “Key” Enthalpies of Formation Revisited. *J. Phys. Chem. A* **2004**, *108*, 9979–9997.
- (72) Ruscic, B.; Bross, D. *Active Thermochemical Tables (ATcT) Values Based on ver. 1.122 of the Thermochemical Network*; Argonne National Laboratory, 2019.
- (73) Campbell, C. T.; Sprowl, L. H.; Árnadóttir, L. Reply to “Comment on ‘Equilibrium Constants and Rate Constants for Adsorbates: Two-Dimensional (2D) Ideal Gas, 2D Ideal Lattice Gas, and Ideal Hindered Translator Models’”. *J. Phys. Chem. C* **2016**, *120*, 20481–20482.
- (74) Johnson, M. S.; Bross, D. H.; Zádor, J. Resolving the Coverage Dependence of Surface Reaction Kinetics with Machine Learning and Automated Quantum. Chemistry Workflows. *J. Phys. Chem. C* **2025**, *129*, 3469–3482.
- (75) Henkelman, G.; Jónsson, H. A dimer method for finding saddle points on high dimensional potential surfaces using only first derivatives. *J. Chem. Phys.* **1999**, *111*, 7010–7022.
- (76) Johnson, M. S.; Pang, H.-W.; Doner, A. C.; Green, W. H.; Zador, J. PySIDT: Subgraph Isomorphic Decision Trees for Molecular Property Prediction. *ChemRxiv* **2024**.
- (77) Zhang, D.; Bi, H.; Dai, F.-Z.; Jiang, W.; Liu, X.; Zhang, L.; Wang, H. Pretraining of attention-based deep learning potential model for molecular simulation. *npj Comput. Mater.* **2024**, *10*, 94.
- (78) Studt, F.; Abild-Pedersen, F.; Varley, J. B.; Nørskov, J. K. CO and CO₂ Hydrogenation to Methanol Calculated Using the BEEF-vdW Functional. *Catal. Lett.* **2013**, *143*, 71–73.
- (79) Göltl, F.; Murray, E. A.; Tacey, S. A.; Rangarajan, S.; Mavrikakis, M. Comparing the performance of density functionals in describing the adsorption of atoms and small molecules on Ni(111). *Surf. Sci.* **2020**, *700*, 121675.
- (80) Wellendorff, J.; Silbaugh, T. L.; Garcia-Pintos, D.; Nørskov, J. K.; Bligaard, T.; Studt, F.; Campbell, C. T. A benchmark database for adsorption bond energies to transition metal surfaces and comparison to selected DFT functionals. *Surf. Sci.* **2015**, *640*, 36–44. Reactivity Concepts at Surfaces: Coupling Theory with Experiment